Report on NIST-ATIS WSTS'05
(Workshop on Synchronization in Telecom Systems)
Broomfield, Colorado – 10-12 May 2005

http://tf.nist.gov/seminars/ATIS.2005.html

Executive Summary: Views from the Workshop

- Major problems may appear for developing telecom due to the following recipe:
  1. Billions of "$£€ are being spent on third generation wireless.
  2. The Next Generation Network (NGN), the shift from synchronous transmission to packet networks, will be a shift as significant as the change from black & white to color TV.
  3. Synchronization is as essential to network elements as power and space.
  4. The sync in the existing network is well engineered, making it transparent. It only becomes visible when there's a failure, which is rare.
  5. Some engineers, many familiar with Internet Protocol (IP), think sync will not be necessary, or perhaps minimally required, in the NGN.
  6. Existing products for transporting sync in the NGN do not meet required specs.
  7. Future sync may require "time-of-day" as well as frequency synchronization.

- Sync engineers are recognizing their responsibility to communicate the importance of sync to senior management before there are failures. They are trying to clearly link sync issues to telecom revenues.

- Good sync requires redundancy. There are many choices for sync sources, and more coming.

- Sync impacts telecom by 1) ensuring within the network that traffic is neither lost nor delayed, and 2) enabling services to customers at the edges of the network.

- With wireless telecom, TV commercials have gone from the paradigm of hearing a pin drop ("That was a pin?") to "Can you hear me now?" Quality of Experience (QoE) is a metric that directly ties to customer satisfaction. Sync for wireless contributes significantly to QoE, reducing dropped calls on hand-offs and improving speech quality.

- Sync management systems can reduce troubles due to sync problems.

- Much of the core network is now based on SONET/SDH equipment, which delivers both data and sync. Many customers depend on this for equipment such as T1/E1 lines. The transition to NGN will require maintaining this backward compatibility.

Report on NIST-ATIS WSTS'05
(Workshop on Synchronization in Telecom Systems)
Broomfield, Colorado – 10-12 May 2005

http://tf.nist.gov/seminars/ATIS.2005.html

WSTS: a Workshop on Sync in Telecom since 1992

The three-day NIST-ATIS Workshop on Synchronization in Telecommunication Systems (WSTS '05) was conducted in Broomfield, Colorado, May 10-12, 2005. This industry-neutral, technology workshop has occurred annually since 1992.There were 81 participants, including 35 speakers from 19 companies. Individuals from component manufacturers, equipment vendors, network operators, power companies, government organizations, the department of defense, industry analysts, and other diverse fields congregated at the Omni Interlocken Resort for three days of talks, panel discussions, and Q&A sessions. This year's attendees included folks from Europe, North and South America, and Hawaii.

The original focus was traditional wireline telecommunication synchronization requirements in North America. The workshop has grown in step with the expanding communications industry to include diverse market segments and attracts participants from all over the world. The workshop offered the opportunity to mix with peers and experts in all areas of telecommunications network timing and synchronization and to meet network operators, strategists, design engineers, system architects, and synchronization planners from many sectors including wireline, wireless, enterprise and utilities.

Timing is critical for the reliable transmission of voice, video and data in any digital communications network. The formal talks ranged from tutorial in nature, covering the basics of timing and synchronization, to advanced presentations addressing forward-looking areas. The interactive nature of the workshop allowed participants to address industry experts with their specific challenges. Questions and answers, comments, observations, opinions, and anecdotes were numerous and informative.

The three-day agenda was organized in a logical flow from "Traditional" (Day 1) to "Transitional" (Day 2) to "Next Generation Networks" (Day 3). Recognizing time limitations

for presentations, extensive versions of presentations were provided to all attendees in hard copy, while the presentations focused on key features.

Day 1: "Traditional"

Session A. *The Importance of Sync and Timing*

A1. Technology Viewpoint – Ken Biholar, (Alcatel, ATIS)
   i. Standards Overview
   ii. The Role of Time and Frequency in Communications

A2. PANEL DISCUSSION on "The Business of Sync", Panel Members: Charles Curry (Chronos), John Yuzdepski (Symmetricom), Mark Jones (Sprint), Jim Harmon (Qwest), Mike Gilson (BT)

Day 1 opened addressing the question of "why are we here?" Why do we need sync? The first session was kicked off by Alcatel's Ken Biholar who is also Chairman of the ATIS standards body OPTXS. Ken is an industry veteran who has played an important role in defining standards for the telecommunications industry for, as he puts it, "longer than I care to admit". He provided an overview of prevailing standards and provided a succinct explanation of the vital role synchronization plays in telecommunications.

Following Ken's talk was a Panel Discussion, *The Business of Sync*, chaired by Charles Curry (Chronos, UK). Panel members included senior executives from equipment manufacturing companies and Service Providers. The panel came into being from a recognition that sync has poor visibility to telecom management. The traditional SONET/SDH systems have synchronization engineered well enough that failures of sync equipment produce few noticeable effects. This may result in sync being under-funded, and even considered not necessary.

As Charles described in his opening remarks, engineers have to "close the ring", demonstrating that synchronization is a key enabler of Quality-of-Service (QoS), QoS is what applications need, and it is the applications from which revenue and profit arise. Likewise engineers need to close the ring by explaining that the impact of bad sync is bad QoS, which in turn leads to unhappy customers and loss of revenue and profits. One panel member provided a useful example where a report on a field trial related to synchronization of wireless base-stations could be condensed, for presentation to a senior executive, into one chart that showed that the number of dropped calls reduced dramatically, improving the *minutes-of-use* metric by 25%, language the executive could appreciate immediately.

The panel explained that building a business case for sync is somewhat more difficult than, for example, building a business case for buying a new switch. Some of the reasons and related observations:

a) It is difficult to identify any increase in revenue that can be attributed to an expenditure on sync equipment. Sync is not usually associated with a service that can be charged for.
b) While sync in and of itself does not bring in revenue, it is an enabler of other services. Consequently a business case has to be based in terms of loss of revenue associated with the lack of sync. In that sense sync is somewhat like insurance and is associated with risk avoidance. This is a more difficult message to convey to senior management.
c) Sync, when working correctly, is "invisible". This is the opposite of "the squeaky wheel gets the grease". Upper management becomes aware of sync only when there is a failure.

Over the years sync has been well engineered and consequently may have received less management attention than it deserves.

d) The capital expenditure on sync is a very tiny fraction of the total expenditure on network equipment. Such "rounding errors" are often ignored by CEOs and CFOs.

e) Sync is an arcane and esoteric field. To increase the level of awareness of sync in the sync community must develop a good "elevator pitch".
The message that sync is an essential infrastructure component, much like space (to mount equipment in racks) or power and ground, has not been adequately conveyed. In a later presentation, it was shown that when translated into a common metric such as "dollars per tap", sync has greater value than power. Yet, expenditures related to power are easily justified by comparison with sync related expenditures.

Session B. *Requirements, Metrics, and Measurements*

B1. Statistics (Metrics) related to Sync and Timing – Marc Weiss (NIST)
B2. SONET/SDH fundamentals – Stephan Bedrosian (Agere)
B3. Sync Performance: Monitoring and Analysis – Lee Cosart (Symmetricom, OPTXS_SYNC)

The second session was tutorial in nature. Dr. Marc Weiss of NIST, a principal organizer of the Workshop since its inception in 1992, provided an overview of the *Statistics(Metrics) related to Sync and Timing*. A version of this presentation has been a staple at every workshop since 1992, educational for newcomers and veterans alike. This was followed by a presentation on *SONET/SDH fundamentals.* The preponderance of core network traffic is carried over SONET/SDH facilities today. This optical network technology will be relevant for the foreseeable future. The third presentation of the session, *Sync Performance: Monitoring and Analysis*, described methodology, equipment, and analysis software tools available for verifying that network elements are indeed functioning as they should be.

Session C. *Primary Reference Sources*

C1. GPS in Telecommunications – Marc Weiss (NIST)
C2. GPS policy and future GPS developments – Hank Skalski (US-DOT)
C3. Galileo – presented by Marc Weiss (NIST)
C4. LORAN as a PRS – Linn Roth (Locus), Tom Celano (TSC)
C5. Cesium as a PRS – Kishan Shenoi (Symmetricom)
C6. Passive Hydrogen Masers for Telecom Applications – Olie Mancini (FEI)

The talk entitled *GPS Policy and future GPS developments,* presented by a representative of the US Department of Transportation, explained the recent policy changes. Of special interest was the elevation of the status of GPS and the reiteration of US Government policy to continually ensure the availability of state-of-the art space based systems supporting a variety of applications. It was especially heartening to the audience that timing services were recognized as being of utmost importance.

We can see that there are a number of choices for PRS clocks. It was helpful to hear the details of each to assist in understanding pros and cons. Good sync in telecom relies on variety and redundancy of sources.

Session D. *Sync in Traditional SONET/SDH/TDM Networks*

D1. Synchronization in a Global Network – Chuck Norman (Sprint)
D2. Qwest's Network – Dennis Coleman (Qwest)
D3. Timing Solutions for Synchronous Telecommunications Equipment – Michael Rupert and Maamoun Seido (Zarlink)

The talks in this last session of Day 1 had two from service providers and the third from a component manufacturer. Service Providers are at the "top of the food chain." They deploy equipment made by network element manufacturers who source chips from semiconductor manufacturers. This session was beneficial to the entire sync community.

As Chuck Norman of Sprint put it, "It's all about choices". Sprint was an early adopter of SONET and has a global SONET/SDH network in place. The original deployment of the North American network in the mid 1980s was publicized by the classic television commercial involving a pin drop with the punch line "*That was a pin?*" This commercial epitomized the quality of experience paradigm for the time. The sync network supporting this extensive network was designed from the ground up to be robust and to require the minimal amount of human intervention. As Chuck explained, Sprint's choice was to make the network as simple as possible and follow a simple set of rules consistently. The emphasis, he said, was on *simplicity* and *consistency*.

The equipment vendors need to be able to support the wide range of choices made by service providers. Likewise, component manufacturers have to able to support the wide range of architectures chosen for the network elements. One point was clear that in spite of the move toward IP, the preponderance of traffic is still carried over SONET.

Day 2: "Transitional"

Session E. *Sync for Wireless*

E1. Synchronizing Wireless Networks – Charles Curry (Chronos)
E2. Wireless: Improving Service Quality with Sync – Barry Dropping (Symmetricom)
E3. Sync in UMTS (Wireless Networks) – Manuel Nardelli (Ericsson)
E4. Synchronization Requirements for CDMA Networks - Brian Harms (Qualcomm)
E5. PANEL DISCUSSION on Sync Issues in Wireless – speakers from this section.

The role of sync for all types of (cellular) wireless networks, including GSM, UMTS, and CDMA was brought out. One talk related synchronization to revenue enhancement, providing

results of a field trial that showed that dropped calls and customer churn were significantly reduced when the level of synchronization was enhanced from "the bare minimum". The fundamental need for both time and frequency in CDMA networks was explained in the presentation from Qualcomm. The need for both time and frequency in TDD implementations of 3G wireless was also explained.

One of the primary topics addressed in the panel discussion was the universal 50 ppb (minimum) requirement for frequency accuracy of the base station clock. Several people in the audience questioned whether this was adequate. This contrasted with the reported view of engineers from the packet world who may think that 50 ppb is too much. The audience and the panel considered it possible that a more stringent requirement would move the prevailing "Can you hear me now" paradigm back to "That was a pin?"

Clearly, this reflected a choice: coverage versus quality. The animated discussion around this made clear that the audience had a "quality-of-experience" mindset and it was generally believed that whereas coverage may be "more important", speech quality in wireless networks could be improved and that sync is one such quality-improvement enabler. Considering mobility is a principal feature of wireless networks, reliable call handoffs are indeed very important and sync is therefore a coverage-improvement enabler as well.

Session F. *Management of Sync Networks*

F1 Network Quality Assurance—Element Management Systems (EMS) for Network Clocks – Ron Evans or Clark Woodward (FEI)
F2. Network Visibility and Control -- Dilip Dhanda (Symmetricom)

Although there were only two presentations in this session, it was clearly communicated that managing networks, especially sync networks was crucial. With the general trend in the industry towards "lights out operation", the need for remotely managed network elements is well understood. For sync networks it is especially important because sync problems such as timing loops can be pernicious and difficult to diagnose. Sync management can make success possible in an environment where only a handful of experts have to monitor and maintain networks, diagnosing and fixing sync problems from a remote location.

Session G. *Emerging Topics in Synchronization and Timing*

G1. Chip-Scale Atomic Clocks – Svenja Knappe (NIST)
G2. Two-Way Time Transfer – Marc Weiss (NIST)
G3. Two-Way Framed Packet Timing – George Zampetti (Symmetricom)
G4. Service Level Agreements for Synchronization Delivery – Ian Wright (Chronos)

The presentation on *Chip-Scale Atomic Clocks* from NIST exposed the audience to the concept of miniaturization of atomic clocks. While still in the development phase, such miniature atomic clocks have the promise to displace high-end ovenized quartz oscillators in scenarios where power consumption is paramount, such as in portable battlefield radios.

There were two presentations on the subject of two-way time-transfer. This technology can distinguish itself from the prevalent (legacy) approach of "send-and-pray" whereby a frequency reference is delivered from the master to the slave (server to the client) and the communication is generally one-way. The master has no feedback from the slave. Two-way time transfer methods involve a full-duplex communication between server and client and can deliver a time reference explicitly and a frequency reference implicitly. The two-way nature allows the master to monitor the performance of the slave.

However, it was noted that the prevalent standards, NTP and IEEE-1588, are both "client centric" whereby the client does the majority of the clock algorithm signal processing and the standards are not explicit in the role the server plays in terms of monitoring and performance management. The accuracy, or lack thereof, attainable by various methods operating at different protocol layers was discussed. Generally speaking, *Layer-3* methods, such as NTP tend to be "in-band" and "routable", providing the greatest flexibility, but are the least accurate in terms of synchronization transfer and performance is "statistical". *Layer-1* methods are usually point-to-point, deterministic, and "out-of-band." They provide the highest level of accuracy. *Layer-2* methods, such as IEEE-1588, can be viewed as somewhere in between. Loosely speaking, the level of accuracy is milliseconds, microseconds, and nanoseconds, as we compare *Layer-3*, *Layer-2*, and *Layer-1* two-way time-transfer methods.

The session was rounded out by the talk *Service Level Agreements for Synchronization Delivery*. This talk showed that synchronization could be positioned as a service, allowing synchronization to be a revenue generator. Such a service requires a methodology to determine when it is functioning and when it is not. This is a fundamental aspect of an SLA that distinguishes it from a standard. Generally speaking, a standard describes performance that must be achieved 100% of the time. The MTIE metric, for example, captures deviations "permanently". An SLA requires some softer thresholds that are geared towards determining whether the service was "fit for purpose" X% of the time or better, with $X < 100$.

Session H. *Preparing for Next Generation Networks and Services*

H1. VoIP Quality of Experience – Kishan Shenoi (Symmetricom)
H2. Evolving your Network – Jim Harmon (Qwest)
H3. Sync in the Evolving Network -- Jim Olsen (Symmetricom)
H4. Sync Transport over Access Technologies – Ian Wright (Chronos)
H5. IEEE-1588 An emerging protocol for delivery of sync – Pat Diamond (Semtech)
H6. Synchronization and the Optical Transport Network – Dominik Schneuwly (Oscilloquartz)

Whereas it is generally accepted that Next Generation Networks are going to follow the packet-switching paradigm, there are numerous services provided today that will have to be continued regardless of whether the network is circuit-switched ("legacy") or "IP". It was pointed out that the legacy transport network ("TDM") is well synchronized and services can "piggy-back" upon it; if and when inter-machine links migrate from SONET (synchronized) to GigE (asynchronous)

the sync reference chain is broken and some services may get "stranded". Presentations H2-H5 addressed methodologies to prevent just this sort of situation.

It was clear from the presentations that the morphing of the network from circuit to packet would be evolutionary and not "fork-lifted". This however has some unique challenges. It was heart-warming to see that service providers, equipment manufacturers, and component vendors in the audience were cognizant of this fact and were prepared to address the problem. On the downside, as one participant observed, recognition of the importance of sync was often limited to just the "sync-cognizant" personnel within the service provider community. Further, this group was shrinking in size as well as "clout". It was clearly a call to action to educate the communication industry at large that the notion "IP does not need synchronization" is shortsighted at best and flat-out wrong when applied to certain services. This is particularly true of services that are real-time in nature.

One talk, *VoIP Quality of Experience*, provided an overview of the manner in which speech quality could be quantified using the "R-value" scale and what the notion of "acceptable quality" would be on this scale. It was noted that for purely voice applications synchronization was less of a factor than delay, echo, compression-loss, and comfort noise. For voice-band services, such as between facsimile machines, sync is indeed important. A suggestion for the maximum frequency offset between the A/D and D/A converters at the two ends of the circuit was provided. An audience member did point out that when and if "relay" methods such as T.38 were deployed in a ubiquitous manner, this requirement could be relaxed.

The session was rounded out by the talk *Synchronization and the Optical Transport Network* (OTN) that provided a description of the emerging technologies in optical networking and the manner in which synchronization is approached.

Day 3: "Next Generation Networks"

Session I. *Sync Requirements for Packet Networks*

I1. Synchronization Issues in Next Generation Networks – Dominik Schneuwly (Oscilloquartz)
I2. PDH over SONET and Packet -- Ravi Subrahmanyan (AMCC)
I3. Timing & Synchronization over Packet Networks -- Jeremy Lewis (Zarlink)
I4. Achieving Circuit Switched QoS in IP/MPLS Networks – Kishan Shenoi (Symmetricom)
I5. Packets Are The Answer? - Synchronisation in the converging world – Mike Gilson (BT)
I6. A Quantitative Study of Timing over Packet Networks – Michel Ouellette (Nortel)
I7. Standards Activities in Synchronization over Packet Networks – Silvana Rodrigues (Zarlink)

It was clear in the design that this third day was saved for the most controversial and newly developing topic: the impact of synchronization, or lack thereof, in Next Generation Networks. This topic seems to have been the hottest, motivating many of the attendees in the first place to come to the Workshop.

The first presentation, *Synchronization Issues in Next Generation Networks*, set the stage for the remaining talks. Here the principles of NGN were explained and the manner in which sync issues could arise was put forward. It was clear that a fundamental premise of NGN was that the transformation from circuit-switched to packet-switched architectures would be invisible to the end-user. NGN would support legacy services with the same QoE as legacy networks and be a flexible platform on which new services could be created.

One of the recurring themes in the presentations was the notion of circuit emulation for transporting legacy constant-bit-rate services, such as T1/E1, over a packet network. The phrase "18 μsec wander limit" featured quite extensively. The concentration on T1 (E1) circuit emulation was fortuitous, more happenstance than by design. Nevertheless, this struck a chord with the audience, considering the vast deployed base of customer-premises equipment that interfaces with the network at the T1 (E1) level and the large number of *private-line* T1 circuits in operation today.

The Nortel presentation in particular was an eye-opener. One of the recurring refrains within the IP community is the possibility of using *adaptive clock recovery* for supporting T1 circuit emulation. Fully three out of the four, ostensibly production quality, units failed to meet the *synchronization wander mask* of G.824 even in a "back-to-back" configuration where the devices were connected by a crossover cable. Even with the limited amount of testing, the presenter expressed uncertainty whether the *traffic wander mask* of G.824 could be met reliably under all loading conditions in a switched, much less routed, environment.

The last part of the session, and the last event of the workshop, was a Panel Discussion, *Sync in Packet*, with the panel comprised of the speakers of Session I.
The panel discussion was kicked off by posing the following question:

Does synchronization play any role in Next Generation networks – networks based on packet-switching methods (asynchronous) and asynchronous transport methods (e.g. Ethernet)?
– Can the network be constructed using free-running clocks in network elements?
Each member of the panel presented a short exposition as to whether he/she believed sync did or did not play a role and why. Following the panel member presentations, the audience was invited to participate with questions, observations, and opinions. The audience response was quite animated. Some of the highlights:

1. A panelist indicated that voice-band services, particularly facsimile transmission, could be provided even if the end-points were not well synchronized. The rationale was that fax transmissions could be replaced by e-mail attachments or by using "relay" methods such as T.38. The counter to this was not technical but legal. Fax transmissions constitute legal documents but scanned documents, even in PDF (or similar) format, are questionable. The legality of fax-relay was not addressed.
2. A multi-decade veteran of the telecommunications industry, sketched out his experience with sync over some telecom history. When the PSTN morphed from analog to digital it was learned, the hard way, that synchronization was required. The experience led to the establishment of sync-related standards and the development of sync networks. Along came ATM, Asynchronous Transfer Mode, with the emphasis on the "asynchronous",

that trumpeted the disestablishment of the sync network and the ability of network elements to "free-run". Following a painful and expensive learning experience, service providers recognized that sync was indeed still necessary. We are now undergoing another sea change, from circuit-switched to packet-switched and the "sync is irrelevant" trumpets are blaring again. He raised the rhetorical question whether history was repeating itself.

3. One panel member raised the point that it was not whether synchronization was relevant but, rather, in what form would synchronization play a role in Next Generation Networks. In particular, he submitted that the traditional notion of *frequency synchronization* (i.e. syntonization) might be less important than *time synchronization* where all network elements have an accurate copy of UTC-traceable *time-of-day*.
4. The notion that *adaptive clock recovery* could support legacy services (e.g. T1 circuit emulation) was discounted almost universally, both by panelists and audience participants. The general argument against adaptive clock recovery was that it was not deterministic enough. There seemed to be a number of examples that network quality timing could not be delivered over the (synthesized) DS1/E1 interface to the terminal equipment that required it. In particular, satisfying G.824 (DS1) and G.823 (E1) could not be guaranteed. It was opined, however, that changing the standards was one approach to allowing adaptive clock recovery methods to become "standards-compliant".
5. Two members of the audience, physicists and self-admittedly not that conversant with telecommunications, suggested that an appropriate model for the emerging packet-switched based NGN was based in Chaos Theory. It was hypothesized that there was possibly one parameter, most likely associated with network loading, that could determine whether the network was usable (for demanding services such as circuit emulation of real-time traffic) or not. In particular, there would be a threshold value for loading such that below the threshold service performance would degrade gracefully with increasing load. However, if the loading increased above the threshold, nothing could be guaranteed.
6. Interesting discussion followed the suggestion that sync was irrelevant for the core (transport) network. The suggestion was that the core would operate "error-free" provided the links between the network elements (routers) were lightly loaded and that techniques such as MPLS were used for traffic engineering. It was observed that the preponderance of traffic, regardless of whether the information was packetized or not, was still carried over SONET/SDH facilities. If the network elements were not synchronized then the transport performance would be adversely affected, even to the point of not being "error-free". A presentation made on Day 1 was referred to where it was explained that whereas 20 ppm clocks in SONET/SDH equipment were "adequate" for providing "keep alive" functionality. The number of pointer movements introduced at this level of frequency offset were probably too excessive to support any reasonable quality of service. It was also pointed out that "error-free" in terms of information bits was necessary but not sufficient when it came to real-time services, where both "bits" and "bit-time" are important. Operating the core asynchronously would not diminish the need for delivering a sync reference to the end-points where the bilateral conversion between circuit and packet is achieved.

The principal conclusion of Day 3 of the Workshop was that synchronization is indeed relevant in Next Generation Networks. This was true for several reasons, not the least of which was that synchronization was a key ingredient to maintain the Quality of Service that end-users expect and have been provided by the existing ("legacy") circuit-switched network.